Harvard-Smithsonian Center for Astrophysics
Precision Astronomy Group
60 Garden St., M/S 63, Cambridge, MA 02138

Date: 1999 July 14 TM99-06
To: Distribution
From: J.F. Chandler and R.D. Reasenberg
Subject: FAME astrometry by centroiding analysis on multiple observations

I. INTRODUCTION

One of the key technological features in the FAME mission is the use of a CCD, both to detect photons and to provide the metric for the positions of target images. A subsequent analysis will derive the star parameters from the pixel-by-pixel photon counts obtained by the CCD. The first stages of that analysis form the topic of this memorandum.

The analysis described here is based on the nominal properties of the instrument outlined in the proposal of August 1998: focal length of 7.5 m, aperture of 0.5 m $\times$ 0.25 m with a central obscuration, and CCD pixels of 15 μm square. (See Reasenberg & Phillips 1998, Phillips & Reasenberg 1998.) More recently, in response to concerns about the accuracy obtainable in centroiding the pixel data, the focal length was increased to 15 m. (The aperture was also expanded to 0.6 m $\times$ 0.25 m and the central obscuration made larger.) By placing more pixels on the diffraction pattern, it was argued, the problem of estimating both the parameters of that diffraction pattern and the position of the centroid would become easier. However, it must be noted that the total number of detected photons per star per observation is unaffected by a focal-length adjustment because the integration time is held constant by a corresponding change in the rotation rate. Thus, the statistical fluctuations in the individual pixel photon counts become more significant as the focal length increases. It will be necessary to examine the centroiding process and its results as a step toward determining the best focal length.

Here, we present the results of a series of numerical experiments demonstrating a proof-of-concept for simultaneous determination of stellar position, magnitude, and temperature, based on a combination of pixel-by-pixel photon counts and four-color photometry, all without an *a priori* star catalog. The technique uses a diffraction code and an externally specified stellar spectrum for calculating the photon count in each pixel or wavelength band, as well as the partial derivatives of those counts with respect to the parameters of interest. The software solves for the parameters by iterative linearized weighted-least-squares estimation. The problem is nonlinear in some respects, but is easily accessible to a combination of "educated" starting values and "steepest descent" estimation (during the first few iterations). We have taken 0.0014 pixel in the star position (or 570 μas in the context of this study) as a requirement for the single-measurement accuracy.

The remainder of this memorandum describes briefly the diffraction code used in this study (Section II) and a starting procedure for the estimator (Section III); it

presents the least-squares results (Section IV), and draws a few conclusions about FAME data reduction (Section V). In particular, we find that the required centroiding accuracy would be possible with the old nominal design. However, this work was limited to main-sequence stars with solar metallicity, and we feel that it needs to be extended to a larger class of targets, including, for example, those in the Gunn-Stryker catalog (Gunn and Stryker 1983). More generally, it should embrace metallicity and stellar surface gravity as free (or perhaps loosely constrained) parameters to be determined as far as possible (to the degree necessary)from the data.

II. SOFTWARE MODEL

The numerical experiments in this study were carried out using a diffraction code developed by J. D. Philips, to be described in a forthcoming Technical Memorandum. The code has been modified in several ways to go beyond sensitivity studies to the actual fitting of the theoretical model to observations, but the core numerical techniques are the same. The software performs a triple integral (by numerical quadrature) over wavelength and the two coordinates in the detector plane to obtain the photon count within each pixel from starlight passing through the specified aperture. The latter two integrals are done quasi-analytically using a 2-dimensional spline fit to a dense grid of points defined in the detector plane. (The required, additional double integral over the aperture is done analytically.) By assumption, the CCD detector is operated in Time Delay and Integrate mode (TDI), such that the accumulating charge in the CCD is shifted in the scan direction at the same rate as the motion of the observed star image due to spacecraft spin. The problem is thus treated as static, and the implicit fourth variable of integration, time, is converted to a simple multiplication of the photon count rates by the time interval. The cross-scan position sensitivity is not considered in this code, and, for simplicity, the star image is assumed to fall exactly on the boundary between two columns of pixels, thereby cutting the computational burden in half through the cross-scan symmetry.

Rather than use a pure black-body spectrum for the diffraction calculations, the code now interpolates the spectrum from a table supplied externally, covering a range of effective temperatures. The tabular interval and the precision for the external spectrum files (2 nm and 5 decimal digits) were chosen such that the errors in photon counts are no more than about one part in 10^5, as shown by direct comparison between the results computed analytically from a black-body spectrum and those obtained with a black-body table. For many of the experiments in this study, the spectral tables were based on a set of stellar spectra supplied by R. L. Kurucz (1999), parametrized by temperature, surface gravity, and metallicity. (See Lejeune et al. 1997, 1998.) Sensitivity to the latter two parameters was not investigated in this study; the full 3-dimensional parameter space was reduced to a single dimension (temperature) by choosing only the solar value of metallicity and interpolating to the main-sequence value of surface gravity for each temperature. Future studies will be needed to determine the effects of the other two parameters on the astrometry.

The aperture is rectangular, 0.5 m $\times$ 0.25 m, and is divided by a central 40% obscuration into two rectangular subapertures. The optical throughput of the device is modeled in the code, but the details need not be discussed in the context of this

memorandum, since the results we are presenting here are independent of any "across-the-board" scaling of the detected pixel-by-pixel photon counts. The spectral response of the device is assumed to be uniform across the entire observing band from 400 nm to 900 nm (and zero outside that band). The measurement uncertainties for the resulting photon counts in the astrometry CCD's are assumed to be due entirely to photon statistics on each pixel. However, the counts from the photometry CCD's (in four colors, evenly spanning the available spectrum from 400 to 900 nm) are taken to be uncertain at the 1% level as a conservative estimate, pending a more careful analysis. (Note that, for a 9th-magnitude star and with the nominal instrument design, the uncertainty in the photometry counts based on photon statistics alone would be typically 0.2%.)

The pixels are taken to be 15 μm square. (For comparision, the full width between first nulls is 18 μm for an unobscured slit 0.5 m wide at the assumed focal length of 7.5 m and the reference wavelength of 600 nm. The detected diffraction pattern is spread out further by blurring in the silicon, but we have not included that effect in this analysis.) Though the pixels are large compared to the desired precision of the the star positions, the results of this study indicate that the high information content in the photon counts for such pixels allows very precise position determinations through straightforward weighted-least-squares fitting, even with the necessity of estimating a star's temperature and brightness (partly) from the same data.

As the spacecraft precesses, the fields of view cross a given star in each of several successive rotations. This sequence of detections is called a "visit." In most cases, a star passing through a field of view is detected at least twice (i.e., in two CCD's) and thus at least four times per rotation. (For a further discussion of visits and their lengths, see Reasenberg 1998). Thus, each star is observed many times over a relatively short time span as the spacecraft precession slowly sweeps the observing plane across the sky. We assume the stellar characteristics are stable over a visit.

The astrometric sensitivity of a single observation varies greatly according to the "pixel phase" (position of the star image relative to the co-moving pixel boundaries). See Phillips (1997) for an illustration of the dependence of sensitivity on pixel phase for a variety of pixel sizes. The data from a single visit naturally cover the whole range of pixel phases in the scan direction, since the accumulated effect of spin-rate variations from one rotation to the next, the non-commensurability of the CCD clock and spacecraft rotation periods, and the unevenness of the CCD chip placement all add a quasi-random component to the star image position on the detector. Thus, we avoid the problem of having to rely only on an "unlucky" pixel phase for determining the star position. (An investigation of statistical fluctuations of the distribution of phases of the observations is beyond the scope of the present study.) Moreover, the biases in the position estimate due to any errors in the temperature or brightness will tend to average out in these circumstances. In order to approximate the condition of thorough coverage of pixel phase, while avoiding exact placement of the star image on either the center or edge of a pixel, this study assumes a collection of seven observations for a star, spaced 1/7 of a pixel apart, starting at a point 1/35 from one edge.

Note that a further improvement is possible. If, at the end of the mission, the spectral description of a given star is inadequate, then there will be a position bias for that star as a function of pixel phase. The general form of such a bias is known.

(Arguments based on symmetry show that the bias is zero both at the pixel edges and at the center and is antisymmetric about the center.) Thus, simple parameter fitting will permit removal of the bias simultaneously with the estimation of parallax and proper motion.

III. PHOTOMETRY-BASED STARTING PROCEDURE

As with any problem in non-linear estimation, it is advantageous to begin with the best available nominal values for the parameters to be determined, and star temperature is the one parameter that most needs such treatment in the present problem. Fortunately, the four-color photometry provides a simple way to estimate the temperature (and magnitude) separately from the astrometric analysis. From the externally supplied stellar spectra, the software computes a set of four normalized counts (one in each of the four photometric bands) for each of 54 temperatures spanning the range 3,000 K to 30,000 K. A similar set of four counts for the target star is generated as part of the pseudo-observations, based on the assumed stellar spectrum. By scanning the table for the best fit to the pseudo-data and then interpolating quadratically about that point, the software can determine the temperature with a formal uncertainty as small as a few degrees for the very cool stars or as large as a few thousand degrees for the very hot stars. The stellar magnitude also follows from this fitting procedure. (We find that magnitude and temperature estimates from the four-color data are correlated, and the correlations are highest at the extremes of star temperature. Thus, these parameters must be estimated together, at least in the starting procedure.)

Figure 1 shows the results of estimating the temperature from noise-free four-color photometry using the same theoretical spectra both for generating the photon counts and for fitting the temperature. As expected, the input temperature is correctly recovered by the fit. It is a simple matter to repeat this exercise using two different and inconsistent models of stellar spectra, to see the effect of incorrectly modeling the spectra in the analysis. Figure 2 shows one such experiment: the normalized 4-color photon counts were computed according to the Kurucz model, but then fitted against a table based on black-body spectra. (This is a deliberately chosen example of fitting with a bad model.) The results agree reasonably well in the temperature range where the spectral peak falls inside the passband of the instrument, but the temparatures increasingly disagree outside that range. Moreover, the disagreements are large compared to the fitting uncertainty (shown via the error bars in Figures 1 and 2, and also shown separately in Figure 3 for the case of fitting with the Kurucz model). In effect, these two models define two different temperature scales for optical photometry. As further investigation shows, even this acute discord over stellar temperature produces only a modest bias in the astrometric results.

IV. RESULTS

The first stage of testing the astrometric fitting was carried out without the benefit of the photometry-based starting procedure described in the previous section. In essence, the objective was to discover in a modest number of test cases whether a grossly wrong starting estimate of the stellar temperature would defeat a brute-force iterative least-squares estimator. The cases examined included a variety of true temperatures and

starting estimates, with differences ranging up to 10,000 K. In several cases with extreme initial temperature offsets, the iterative estimator diverged, but a modified estimator using a "steepest descent" technique (see below) was able to converge even in these cases. These worst cases sometimes required many iterations to reach convergence (as many as 15 in the cases studied). A more exhaustive study might be used to determine the ultimate robustness of the modified estimator against even more extreme starting errors. However, astrometric results for the huge numbers of stars to be studied by FAME demand not just convergence, but rapid convergence. It is clear from these initial studies that some starting procedure, like the photometry-based one described above, is needed for the astrometry. Note that throughout this analysis, the convergence criterion was that each of the parameter adjustments be no more than 10% of the corresponding statistical standard deviation.

The well-known weakness of linearized least-squares estimation for non-linear problems is that the scale of the adjustments can be wildly wrong when the assumed parameter values are too far from the correct values. In some cases, there may be confusing local minima in the sum-squared deviation; even in the absence of those, a non-linear problem can produce parameter adjustments of the wrong size (although of the right sign). In order to prevent this estimation scheme from diverging or oscillating out of control in such cases, we imposed upper limits on the adjustments per iteration of the various parameters. Whenever the least-squares estimator generated adjustments that exceeded one or more of these limits, all of the adjustments were scaled down by a common factor until the limits were obeyed. We found that a limit of 1500 K on the star temperature and 0.1 pixel on the star position prevented divergence without requiring many extra iterations of the estimator, but modest changes to these limits would be equally satisfactory.

The next stage of testing entailed iterative solutions using the starting procedure, with a goal of learning the rate of convergence. The results were both encouraging and disappointing. The estimator converged much sooner than with arbitrary starting temperatures, but the basic difficulty in this analysis is the non-linearity of the parameter set. Even with "good" starting values for the temperature and magnitude, the estimator still lacks a procedure for initializing the star positions. As it stands (with position starting values arbitrarily set in each case to the center of the pixel where the image appears), the procedure requires typically 5-8 iterations to reach convergence. We expect that a simple interpolation from the two highest pixel counts, probably making use of the temperature derived from the photometry, would yield much better starting values for the star positions, but we have not yet developed such an algorithm.

We tried a further modification to the estimator, namely, holding the temperature fixed for a few iterations before setting it free to achieve final convergence. Because the estimation problem is so non-linear, and yet converges to very nearly the correct star positions with the temperature held fixed at a modest offset from the correct value, we hoped that such a procedure would lead to faster convergence. Indeed, in the present configuration of the software, the "inner" loop of the triple integration (i.e., wavelength) is redone when and only when the temperature parameter changes, thus requiring much less computer time for an iteration when the temperature is held fixed. This experiment therefore reduced the time needed to reach convergence, but, unfortunately, not the total

number of iterations. In processing the mission data, we envision the use of sub-optimal estimation techniques which greatly reduce the amount of computer time needed for an iteration, and which do not necessarily bestow any advantage on iterations that keep the temperature fixed. Therefore, we needed to find an even faster method.

One such method is to determine the temperature solely from the four-color photometry and thereby avoid the most non-linear part of the problem altogether. We evaluated the astrometric errors that result from this method by the following procedure. First, we calculated the weighted RMS position bias in the "representative" seven pixel phases for a temperature arbitrarily offset 100 K from the correct value, but with all the other parameters converged. Since the dummy data were noise-free and computed from the same model as was used for fitting, the bias was entirely due to the temperature offset. We chose the offset of 100 K to be within, or at least close to, the linear regime and yet large enough that the results would not be lost in numerical noise. The resulting biases are shown in Figure 4. In the linear approximation, Figure 4 is just 100 times the partial derivative of the position bias with respect to temperature. We combined that partial derivative with the temperature uncertainty from Figure 3 to obtain the RMS position bias due to holding the temperature fixed at the value obtained from the photometry, shown in Figure 5. As described in Section II, the uncertainty assumed for the photometric data was a conservative 1% for each measurement, but the true uncertainties may be much better than that (for $V \leq 9$). With only a modest improvement in the assumed photometry, the curve in Figure 5 would lie below 10^{-3} pixel across the entire range of temperatures 3,000 to 30,000 K.

The same experiment could be performed using a different model of stellar spectra. Figure 6 shows the result of fixing the temperature at 100 K from the "true" temperature, with both the dummy data and the fitting carried out using a black-body model. (By "true" temperature, we mean the effective temperature used in the black-body computations of the dummy data.) The results are so similar to Figure 4 that we dispensed with the rest of the experiment.

We next considered what might happen if an inappropriate model of the stellar spectrum were used for fitting to the data. As Figure 2 shows, the temperature estimates from the photometry can be wildly inaccurate if a black-body model is used for fitting data computed using the Kurucz model. We followed the full least-squares fitting procedure with this same model mismatch: the photometric and astrometric pseudo-data were both computed with Kurucz spectra and then both fitted with black-body spectra. The star position estimates were then compared with the true positions, and the weighted RMS bias was computed for each temperature. The results are shown in Figure 7.

The plot does not cover the full temperature range from 3,000 to 30,000 K because of the large discrepancies between Kurucz and black-body models. Since the tables used by the fitting program cover the same nominal temperature range in both cases, Kurucz temperatures above about 20,000 K and below about 3,500 K are "off scale" when fitted with the black-body tables. In any case, the figure shows that, even with an incorrect spectral model and the resulting large errors in the temperature estimates, this procedure gives astrometric biases under 10^{-3} pixel for all but the coolest stars (for which the black-body model is especially inaccurate).

To continue the investigation of this same mismatch (pseudo-data computed from the Kurucz model, but fitted according to the black-body model), we also examined the bias contribution from using the temperatures derived from the (mismatched) starting procedure without further adjustment in the least-squares estimator. As before, we obtained the sensitivity to a temperature offset by arbitrarily fixing the temperatures and comparing the position biases. In this case, the procedure required an extra step, since the "true" temperatures used in calculating the dummy photon counts would not be matched by least-squares fitting with black-body spectra, and the best-fit position biases would therefore not be zero. We computed the position biases for all seven observations of the star at each temperature using the least-squares estimator for position, temperature, and magnitude, and then substracted these from the corresponding biases obtained from holding the temperatures fixed at 100 K from the best-fit values. We then formed the RMS of these differential biases, and the results are plotted in Figure 8. These are similar to, but significantly smaller than, those shown in the corresponding Figures 4 and 6, the most obvious difference being that the sensitivity falls off here at higher temperatures. Since the sensitivity is smaller, it comes as no surprise that the result of combining this with the uncertainty of the mismatched starting procedure (seen in the error bars of Figure 2) yields a correspondingly smaller contribution to the bias, plotted in Figure 9.

V. DISCUSSION

The results of these experiments fall into two categories: descriptive and prescriptive. In the first category, we find that least-squares estimation procedure obtains correct results for position, temperature, and magnitude from the raw pixel-by-pixel photon counts, and that the procedure works much more reliably when four-color photometry is included. It remains to be demonstrated whether the four-color photometry is usefully sensitive to the secondary parameters of the Kurucz spectral model (metallicity and surface gravity), and, indeed, whether errors in those parameters might result in appreciable astrometric biases. That is left to a further study.

We note that the star temperature and magnitude may vary over time, so that, in principle, each "visit" would require an independent determination of these photometric parameters. In practice, however, variable stars typically follow a repeating cycle that could be modeled in the data reduction. In the same way, the surface gravity for a pulsating star would vary over the star's cycle. However, except in the case of close binary systems with accretion disks or cataclysmic variables, the metallicity should be a constant over the time scale of a FAME mission.

In addition, we have shown that a quick estimate of the stellar temperature, based solely on the photometric data, is good enough for the desired level of astrometric accuracy without further refinement through least-squares estimation (subject to confirmation of the measurement uncertainties in those data).

In the second category, the main conclusion from this study is that the photometric data are important to the astrometric part of the mission, both because of the improved temperature sensitivity and because of the increased speed in processing the astrometric data.

VI. REFERENCES

Gunn, J. E., & Stryker, L. L., 1983, ApJ Supp, 52, 121-153.

Kurucz, R. L. 1999, Private communication

Lejeune, T., Cuisinier, F., & Buser, R., 1997, A&A Supp, 125, 229-246.

Lejeune, T., Cuisinier, F., & Buser, R., 1998, A&A Supp, 130, 65-75.

Phillips, J. D., 1997, presentation at ARI Heidelberg, 8 July 1997, available on the FAME web site or from the author, jphillips@cfa.harvard.edu.

Phillips, J. D., & Reasenberg, R. D. 1998, in SPIE Proceedings 3356, Space Telescopes and Instruments V, ed. P. Y. Bely and J. B. Breckenridge (Bellingham: SPIE), 832.

Reasenberg, R. D., 1998, technical memorandum TM98-07.

Reasenberg, R. D., & Phillips, J. D. 1998, in SPIE Proceedings 3356, Space Telescopes and Instruments V, ed. P. Y. Bely and J. B. Breckenridge (Bellingham: SPIE), 622.

FIGURE CAPTIONS

Figure 1. The results of fitting the Kurucz model to dummy 4-color photometry data calculated from the Kurucz stellar spectra. The error bars show the formal standard deviations of the fits based on the assumption of 1% uncertainties in all of the photometric measurements. The spectral classes are indicated along the bottom of the figure (the letter "M" marking M0, and so on).

Figure 2. The results of fitting the black-body model to dummy 4-color photometry data calculated from the Kurucz stellar spectra. The error bars show the formal standard deviations of the fits based on the same assumed photometric data uncertainties as in Figure 1.

Figure 3. The formal standard deviations from Figure 1 are here plotted on a log scale. These values represent the uncertainties in the starting procedure for temperature in the least-squares parameter estimation.

Figure 4. The weighted RMS of the seven star position biases in each combined fit of seven observations of a single star. The temperature is fixed in each fit at 100 K from the true value. The Kurucz spectral model is used both for calculating the dummy data and for fitting those data.

Figure 5. The results of multiplying the temperature uncertainties from Figure 3 by the astrometric sensitivity implied by Figure 4. These are the formal standard deviations in the star positions due to fixing the temperature at the starting value from the photometry.

Figure 6. The weighted RMS of the seven star position biases in each combined fit of seven observations of a single star. The temperature is fixed in each fit at 100 K from the true value. The black-body model is used both for calculating the dummy data and for fitting those data.

Figure 7. The weighted RMS of the seven star position biases in each combined fit of seven observations of a single star. The temperature is estimated in each fit along with the position and magnitude. The input data are computed from the Kurucz spectra, but the solutions use a black-body model to fit the data.

Figure 8. The weighted RMS of the differences in the seven star position biases between the best-fit case and a case with a temperature offset of 100 K. The input data are computed from the Kurucz spectra, but the solutions use a black-body model to fit the data.

Figure 9. The results of multiplying the temperature uncertainties in the starting procedure by the astrometric sensitivity implied by Figure 8. These are the formal standard deviations in the star positions due only to fixing the temperature at the starting value from the photometry.

DISTRIBUTION:

SAO:

R.W. Babcock FAME web site via Marc Murison
J.F. Chandler
J.D. Phillips
R.D. Reasenberg
I.I. Shapiro

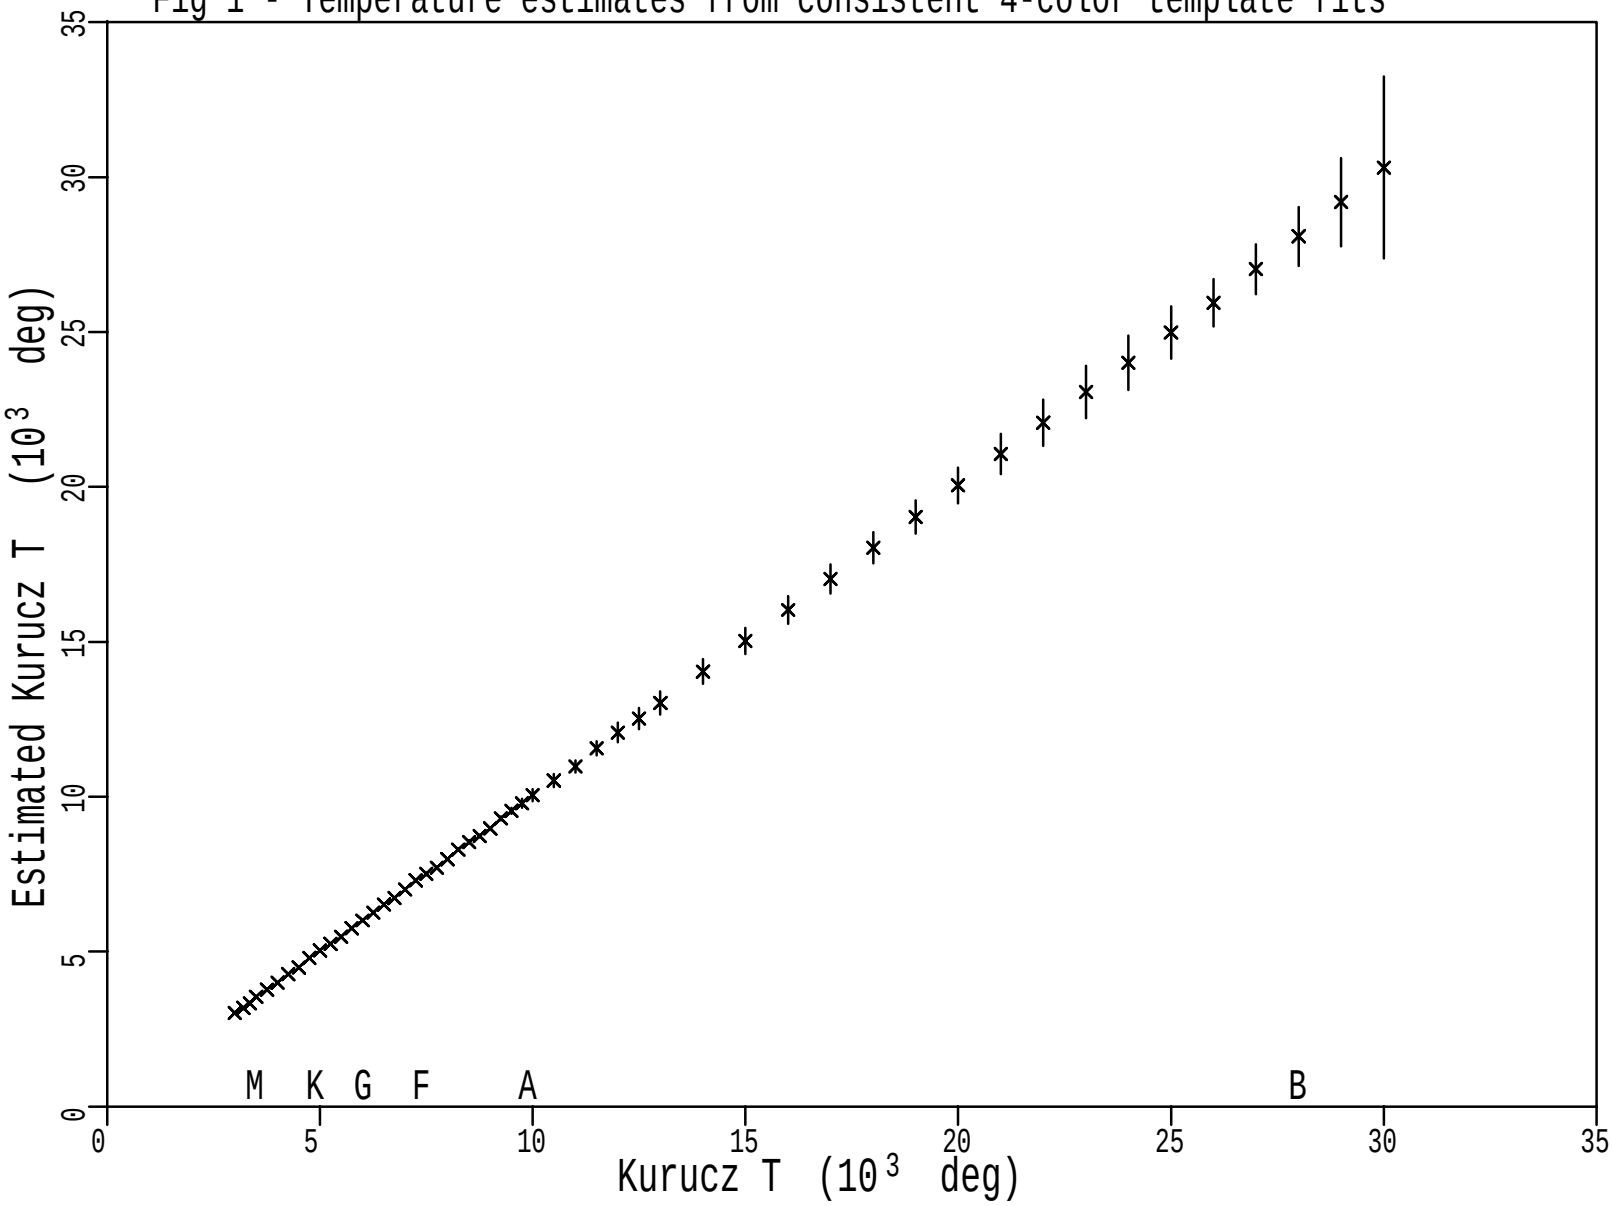
Fig 1 - Temperature estimates from consistent 4-color template fits
Estimated Kurucz T (10^3 deg)
Kurucz T (10^3 deg)
0
5
10
15
20
25
30
35
M
K
G
F
A
B

Fig 2 - Temperature estimates from inconsistent 4-color template fits

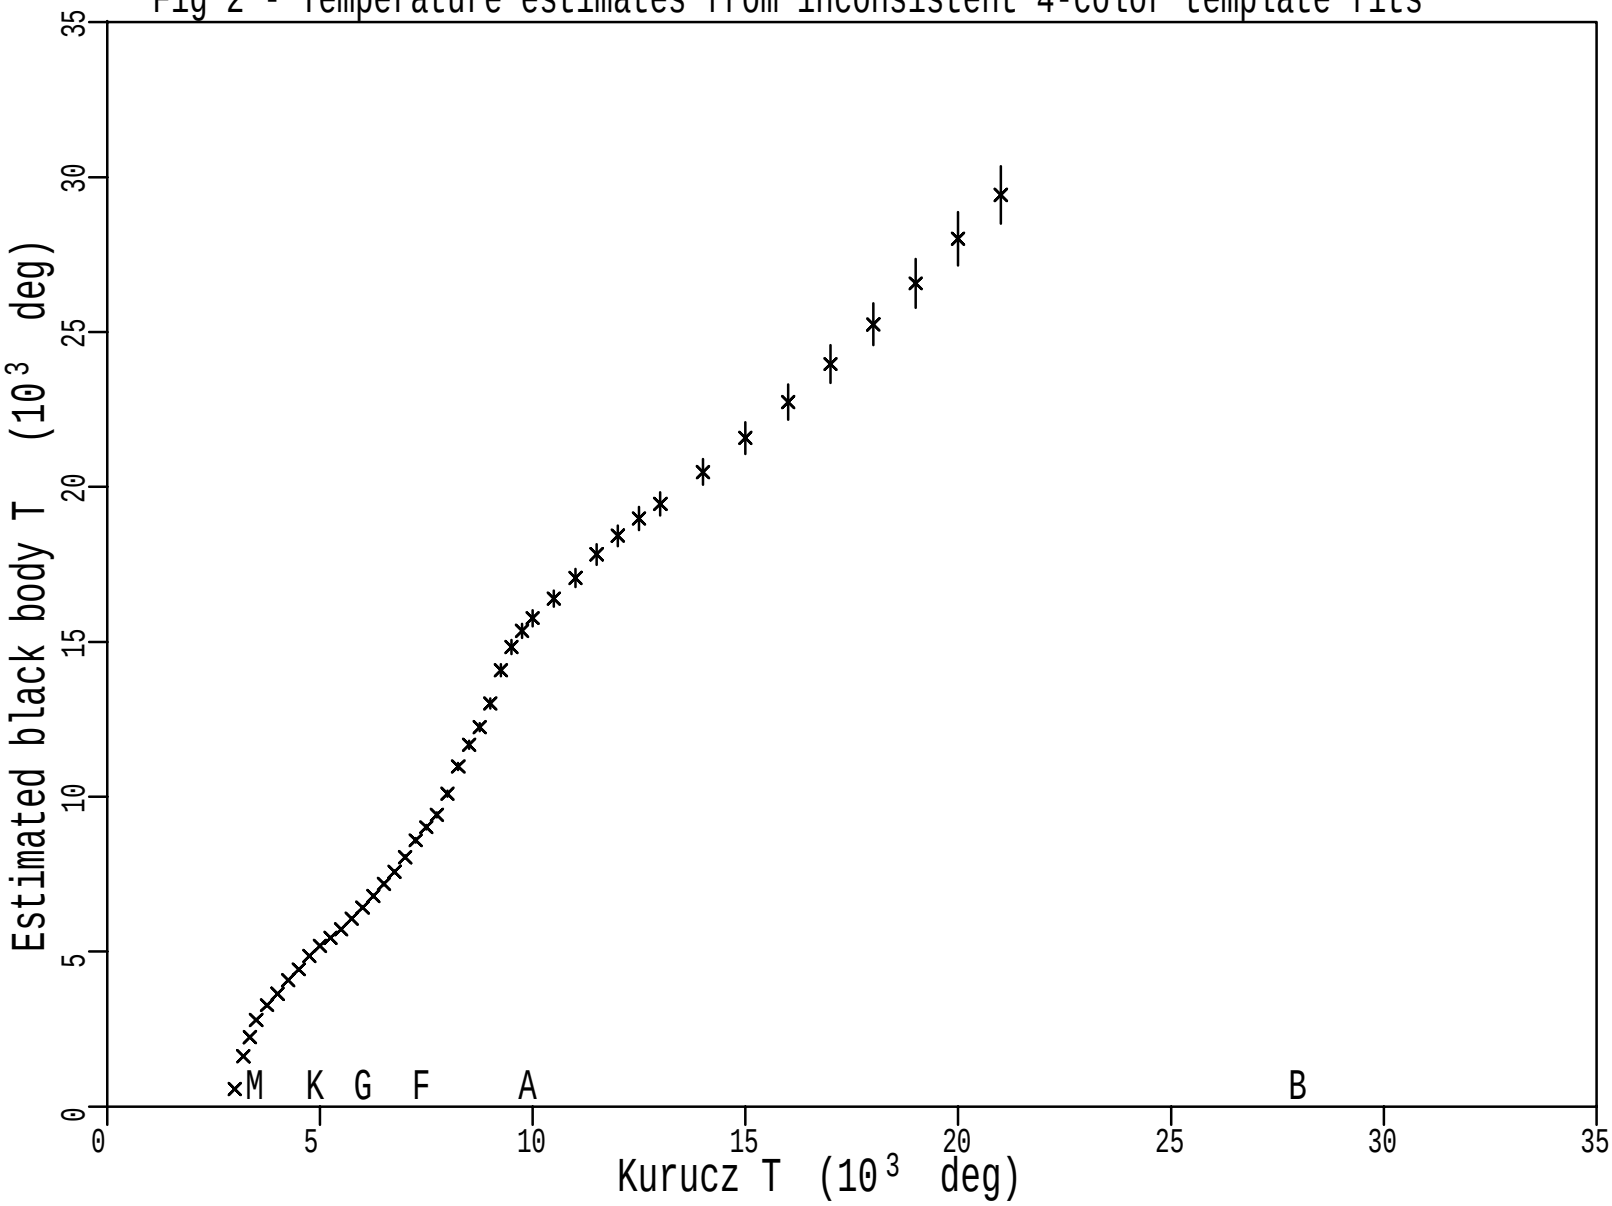

Fig 3 - Temperature uncertainty in starting procedure

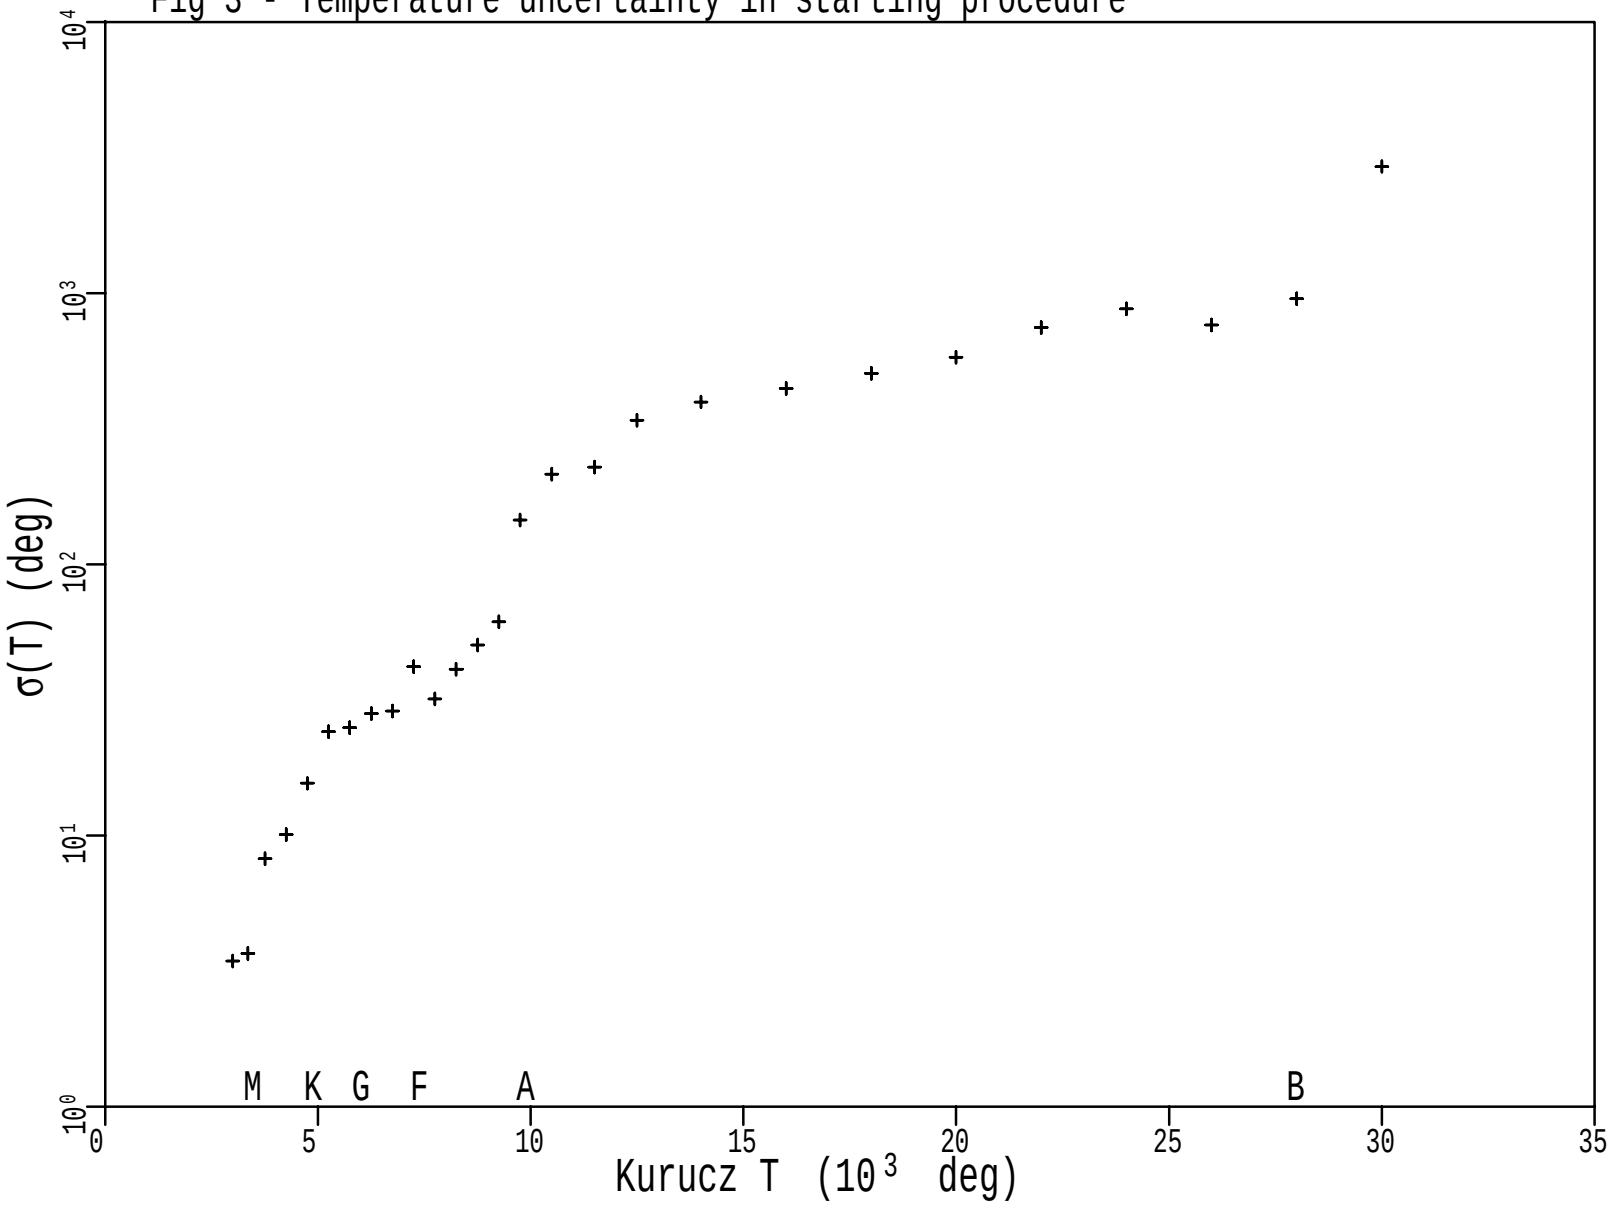

Fig 4 - RMS position bias due to 100 K temperature offset

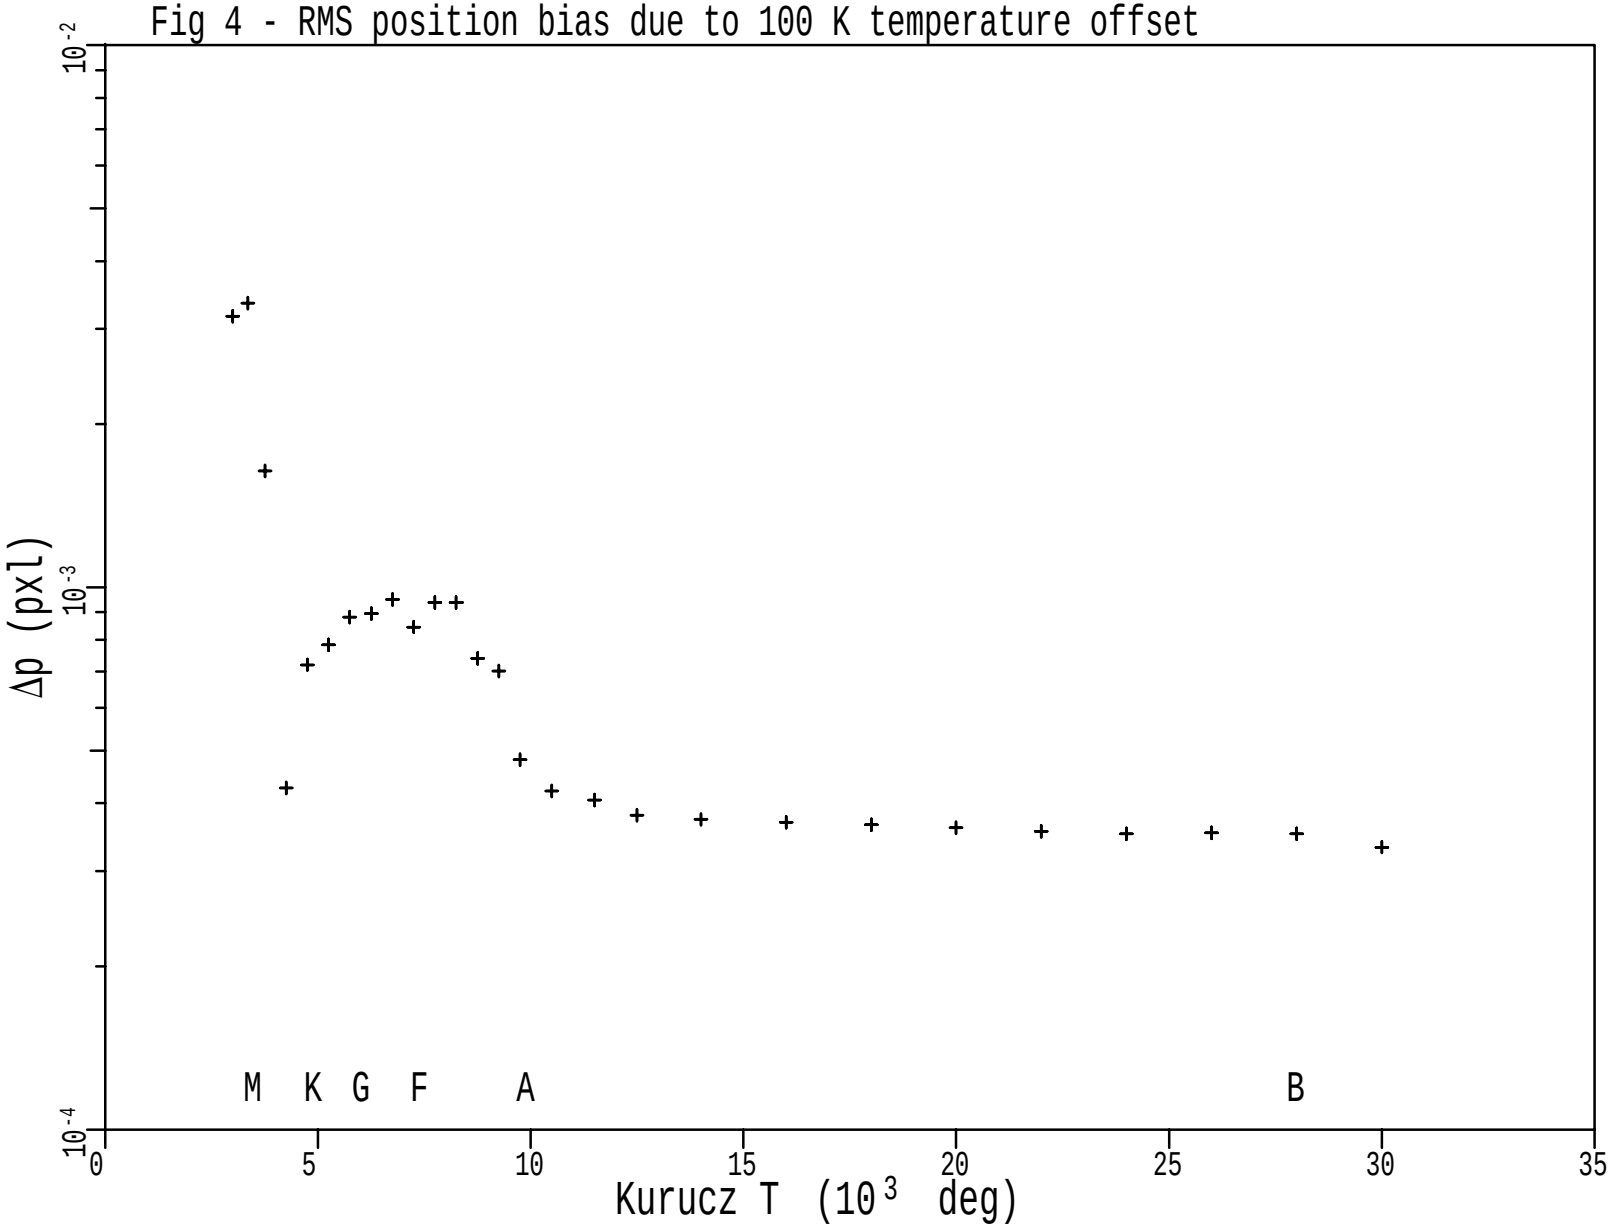

Fig 5 - RMS position uncertainty due to starting procedure

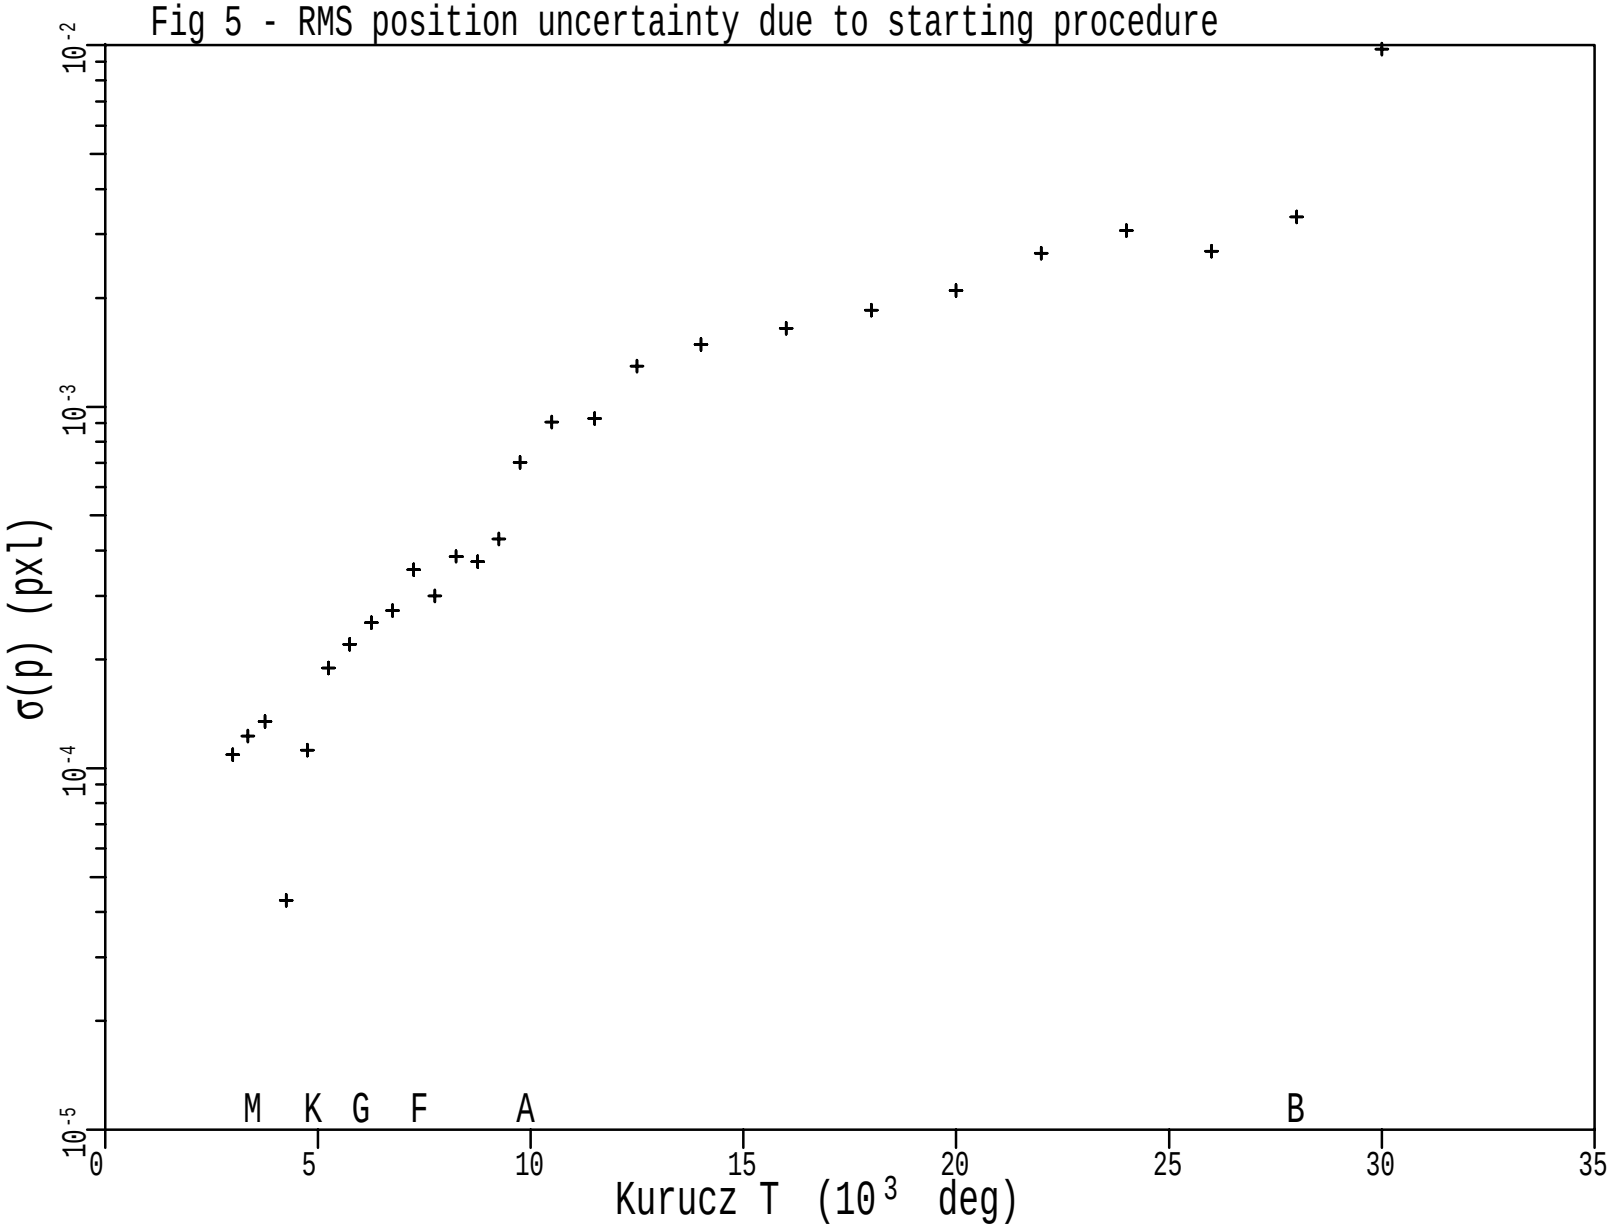

Fig 6 - RMS position bias due to 100 K temperature offset

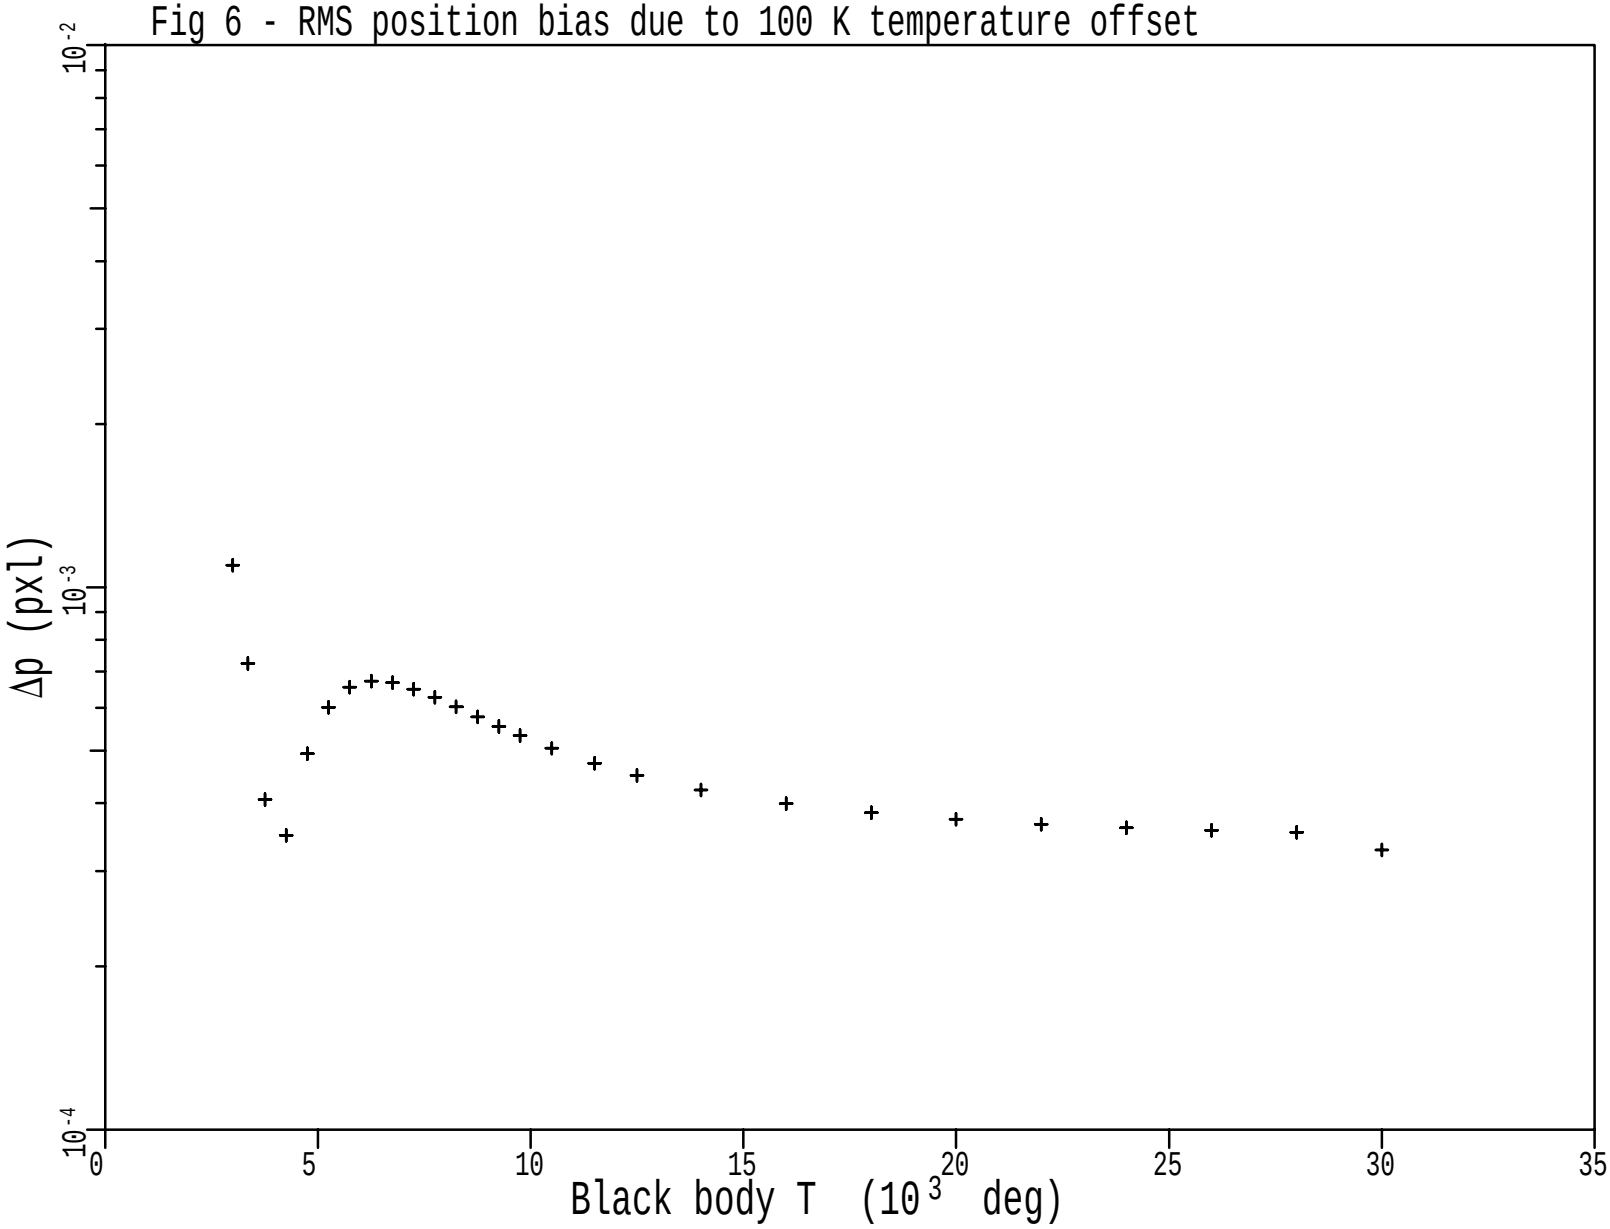

Fig 7 - RMS position bias fitting Kurucz data with black body model

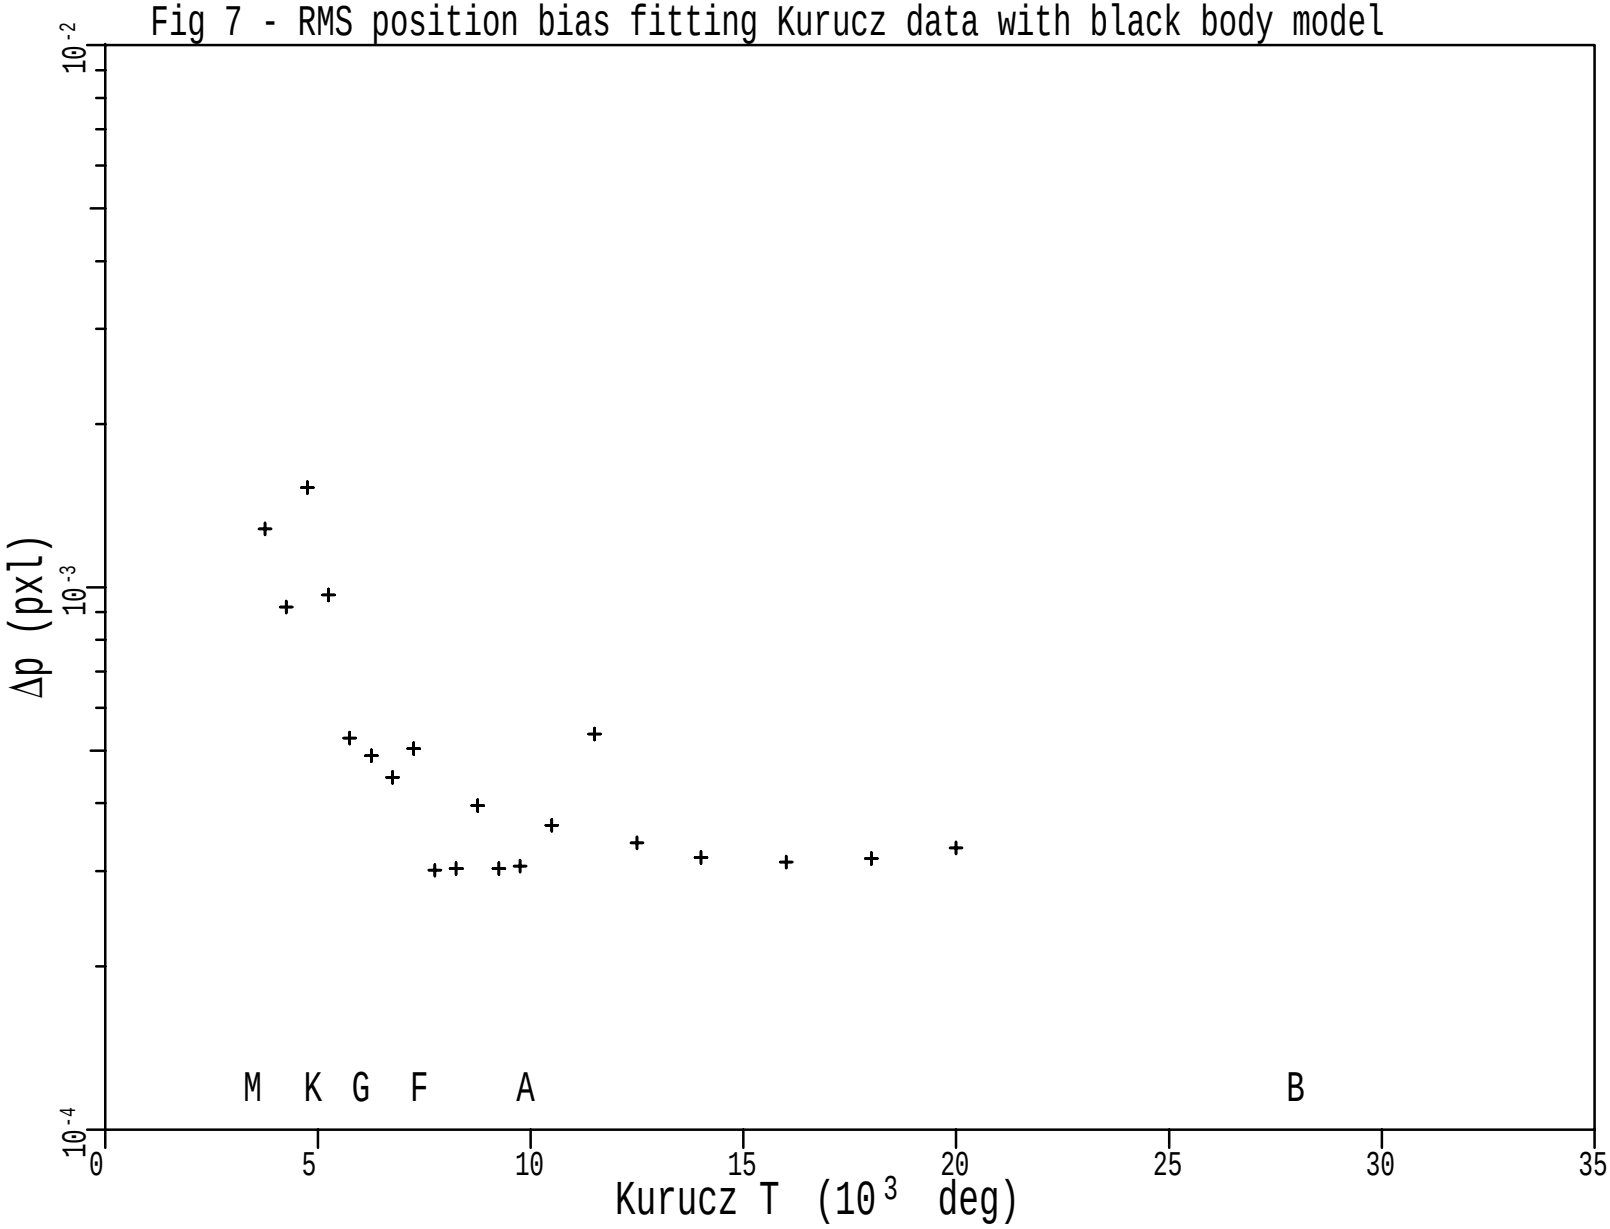

Fig 8 - RMS differential position bias due to 100 K temperature offset

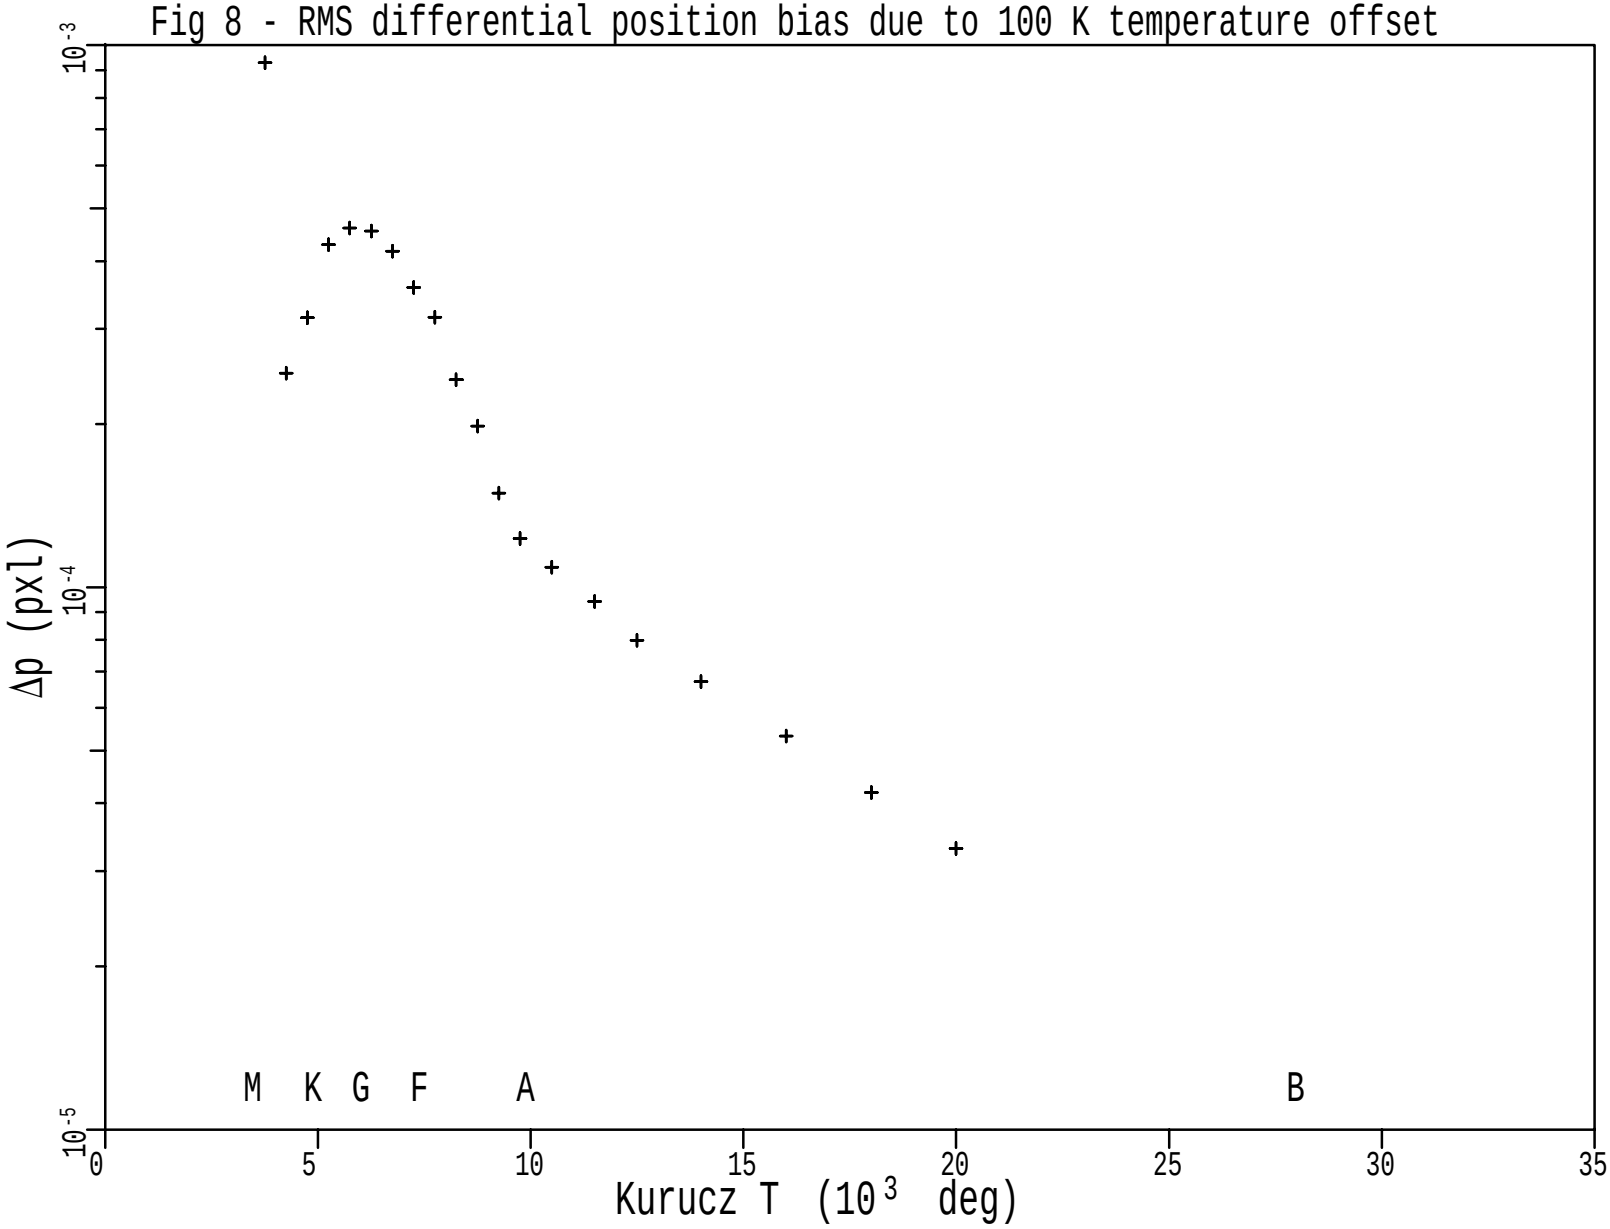

Fig 9 - RMS position uncertainty due to black body starting procedure